

WHAT IS CLAIMED IS:

1 1. A scheduling apparatus for servicing a plurality of data packet flows,
2 each of said data packet flows having specified service requirements and being
3 associated with respective one of a plurality of queues, each of said plurality of
4 queues containing data packets, the distribution of service to said data packet
5 flows being organized in service frames, each service frame offering a number of
6 service opportunities to each of said data packet flows, said number of service
7 opportunities depending on the service requirements of said data packet flows, the
8 scheduling apparatus comprising:

9
10 a first scheduler providing service to a first set of guaranteed-bandwidth
11 (GB) data packet flows during a first subframe of a service frame, said GB data
12 packet flows having service requirements expressed in terms of bandwidth
13 requirements, the duration of the first subframe being adjusted to satisfy the
14 bandwidth requirements of said GB data packet flows; and

15
16 a second scheduler providing service to a second set of best-effort (BE)
17 data packet flows during a second subframe of the service frame, said BE data
18 packet flows having service requirements that are not expressed in terms of
19 guaranteed bandwidth requirements, the duration of the second subframe being
20 the duration of the service frame minus the duration of the first subframe in the
21 same service frame.

1 2. The scheduling apparatus of claim 1, wherein the second scheduler
2 equally allocates service to each data packet flow in said second set of BE data
3 packet flows.

1 3. The scheduling apparatus of claim 1, wherein the first scheduler is a
2 Weighted Round Robin (WRR) scheduler.

1 4. The scheduling apparatus of claim 3, wherein the second scheduler is a
2 WRR scheduler.

1 5. The scheduling apparatus of claim 1, wherein the duration of the service
2 frame is determined to accommodate the sum of the bandwidth requirements of
3 the GB data packet flows in the case where said sum of the bandwidth
4 requirements is equal to said fixed service rate.

1 6. The scheduling apparatus of claim 1, wherein the duration of the service
2 frame is fixed.

1 7. The scheduling apparatus of claim 1, wherein the duration of the first
2 subframe is fixed.

1 8. The scheduling apparatus of claim 1, wherein the duration of the service
2 frame is variable.

1 9. The scheduling apparatus of claim 1, wherein the duration of the first
2 subframe is variable.

1 10. The scheduling apparatus of claim 1, wherein the end of the second
2 subframe is determined when either (1) a predetermined duration of the service
3 frame expires, or (2) there is no remaining backlog in said second set of BE flows,
4 whichever condition occurs first, a backlog being defined as the presence of one
5 or a plurality of data packets in one or a plurality of queues associated with the set
6 of data packet flows being considered.

1 11. The scheduling apparatus of claim 1, wherein the duration of said
2 second subframe is derived from a BE running share ϕ_{BE} , which is defined as the
3 difference between said fixed service rate and the sum of the bandwidth
4 requirements of each of the GB flows that have a backlog in respective ones of
5 said plurality of queues.

1 12. The scheduling apparatus of claim 11, wherein the BE running share
2 ϕ_{BE} receives an increment every time a GB flow becomes idle and a decrement
3 every time a GB flow becomes backlogged.

1 13. The scheduling apparatus of claim 12, wherein said increment and
2 decrement are related to the bandwidth requirement of the GB data packet flow
3 that becomes idle and backlogged, respectively.

1 14. The scheduling apparatus of claim 11, wherein a BE cumulative share
2 Φ_{BE} is set equal to said BE running share ϕ_{BE} at the end of the first subframe, and
3 equal to zero at the end of the service frame.

1 15. The scheduling apparatus of claim 14, wherein the selection of the
2 second scheduler has strict priority over the selection of the first scheduler if the
3 cumulative share Φ_{BE} is not equal to zero.

1 16. The scheduling apparatus of claim 1, wherein said first scheduler
2 maintains a timestamp F_{BE} for the aggregate of BE data packet flows in order to
3 regulate the duration of said second subframe.

1 17. The scheduling apparatus of claim 1, wherein:
2

3 the data packet flows that are backlogged in said first set of GB flows are
4 sorted for service in a Primary WRR Scheduler (PWS) First-In-First-Out (FIFO)
5 queue;

6
7 each data packet flow of said first set of GB flows has a frame flag
8 associated therewith; and

9
10 said first subframe is over when the frame flag of the GB flow at the head
11 of said PWS FIFO queue has different value than a global frame counter GFC .

1 18. The scheduling apparatus of claim 17, wherein the value of said GFC
2 is set equal to the value of said frame flag of said GB flow at the head of said
3 PWS FIFO queue after a difference between the two values is detected.

1 19. The scheduling apparatus of claim 16, wherein the data packet flows
2 that are backlogged in said second set of BE data packet flows are sorted for
3 service in a Secondary WRR Scheduler (SWS) FIFO queue.

1 20. The scheduling apparatus of claim 17, wherein the end of said second
2 subframe is declared every time the timestamp F_{BE} of said aggregate of BE data
3 packet flows exceeds a PWS reference timestamp increment T_Q^{PWS} .

1 21. A method of operating a scheduler for servicing a plurality of data
2 packet flows, each of said data packet flows having specified service requirements
3 and being associated with one of a plurality of queues, each of said plurality of
4 queues containing data packets, the distribution of service to said data packet
5 flows being organized in service frames, each service frame offering a number of
6 service opportunities to each of said data packet flows, said number of service

7 opportunities depending on the service requirements of said data packet flows, the
8 method comprising the steps of:

9

10 providing service to a first set of guaranteed-bandwidth (GB) data packet
11 flows during a first subframe of a service frame, said GB data packet flows
12 having service requirements expressed in terms of bandwidth requirements, the
13 duration of the first subframe being adjusted to satisfy the bandwidth
14 requirements of said GB data packet flows; and

15

16 providing service to a second set of best-effort (BE) flows during a second
17 subframe of the service frame, said BE data packet flows having service
18 requirements that are not expressed in terms of guaranteed bandwidth
19 requirements, the duration of the second subframe being the duration of the
20 service frame minus the duration of the first subframe in the same service frame.